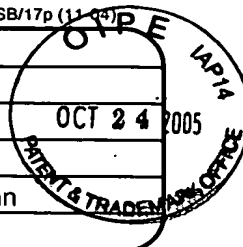


PETITION FEE
Under 37 CFR 1.17(f), (g) & (h)
TRANSMITTAL

(Fees are subject to annual revision)

Send completed form to: Commissioner for Patents
P.O. Box 1450, Alexandria, VA 22313-1450

Application Number	10/767,247
Filing Date	January 28, 2004
First Named Inventor	Watanabe, Naoki
Art Unit	2188
Examiner Name	Mano Padmanabhan
Attorney Docket Number	16869B-077300US



Enclosed is a petition filed under 37 CFR §1.102(d) that requires a processing fee (37 CFR 1.17(f), (g), or (h)). Payment of \$ 130.00 is enclosed.

This form should be included with the above-mentioned petition and faxed or mailed to the Office using the appropriate Mail Stop (e.g., Mail Stop Petition), if applicable. For transmittal of processing fees under 37 CFR 1.17(i), see or PTO/SB/17i.

Payment of Fees (small entity amounts are NOT available for the petition fees)

- ☒ The Commissioner is hereby authorized to charge the following fees to Deposit Account No. 20-1430 :
☒ petition fee under 37 CFR 1.17(f), (g) or (h) ☒ any deficiency of fees and credit of any overpayments
Enclose a duplicative copy of this form for fee processing.

☐ Check in the amount of \$ _____ is enclosed.

☐ Payment by credit card (Form PTO-2038 or equivalent enclosed). Do not provide credit card information on this form.

Petition Fees under 37 CFR 1.17(f): Fee \$400 Fee Code 1462

For petitions filed under:

- § 1.53(e) - to accord a filing date.
- § 1.57(a) - to accord a filing date.
- § 1.182 - for decision on a question not specifically provided for.
- § 1.183 - to suspend the rules.
- § 1.378(e) - for reconsideration of decision on petition refusing to accept delayed payment of maintenance fee in an expired patent.
- § 1.741(b) - to accord a filing date to an application under § 1.740 for extension of a patent term.

Petition Fees under 37 CFR 1.17(g): Fee \$200 Fee Code 1463

For petitions filed under:

- § 1.12 - for access to an assignment record.
- § 1.14 - for access to an application.
- § 1.47 - for filing by other than all the inventors or a person not the inventor.
- § 1.59 - for expungement of information.
- § 1.103(a) - to suspend action in an application.
- § 1.136(b) - for review of a request for extension of time when the provisions of section 1.136(a) are not available.
- § 1.295 - for review of refusal to publish a statutory invention registration.
- § 1.296 - to withdraw a request for publication of a statutory invention registration filed on or after the date the notice of intent to publish issued.
- § 1.377 - for review of decision refusing to accept and record payment of a maintenance fee filed prior to expiration of a patent.
- § 1.550(c) - for patent owner requests for extension of time in ex parte reexamination proceedings.
- § 1.956 - for patent owner requests for extension of time in inter partes reexamination proceedings.
- § 5.12 - for expedited handling of a foreign filing license.
- § 5.15 - for changing the scope of a license.
- § 5.25 - for retroactive license.

Petition Fees under 37 CFR 1.17(h): Fee \$130 Fee Code 1464

For petitions filed under:

- § 1.19(g) - to request documents in a form other than that provided in this part.
- § 1.84 - for accepting color drawings or photographs.
- § 1.91 - for entry of a model or exhibit.
- ☒ § 1.102(d) - to make an application special.
- § 1.138(c) - to expressly abandon an application to avoid publication.
- § 1.313 - to withdraw an application from issue.
- § 1.314 - to defer issuance of a patent.

Chun-Pok Leung
Signature

Chun-Pok Leung

Typed or printed name

October 24, 2005

Date

41,405

Registration No., if applicable

10-26-05

ZPW

PTO/SB/21 (09-04)



TRANSMITTAL FORM

used for all correspondence after initial filing)

Total Number of Pages in This Submission

13

Application Number

10/767,247

Filing Date

January 28, 2004

First Named Inventor

Watanabe, Naoki

Art Unit

2188

Examiner Name

Mano Padmanabhan

Attorney Docket Number

16869B-077300US

ENCLOSURES (Check all that apply)

<input checked="" type="checkbox"/> Petition Fee Transmittal <input type="checkbox"/> Fee Attached <input type="checkbox"/> Amendment/Reply <input type="checkbox"/> After Final <input type="checkbox"/> Affidavits/declaration(s) <input type="checkbox"/> Extension of Time Request <input type="checkbox"/> Express Abandonment Request <input type="checkbox"/> Information Disclosure Statement <input type="checkbox"/> Certified Copy of Priority Document(s) <input type="checkbox"/> Reply to Missing Parts/ Incomplete Application <input type="checkbox"/> Reply to Missing Parts under 37 CFR 1.52 or 1.53	<input type="checkbox"/> Drawing(s) <input type="checkbox"/> Licensing-related Papers <input checked="" type="checkbox"/> Petition to Make Special <input type="checkbox"/> Petition to Convert to a Provisional Application <input type="checkbox"/> Power of Attorney, Revocation Change of Correspondence Address <input type="checkbox"/> Terminal Disclaimer <input type="checkbox"/> Request for Refund <input type="checkbox"/> CD, Number of CD(s) _____ <input type="checkbox"/> Landscape Table on CD	<input type="checkbox"/> After Allowance Communication to TC <input type="checkbox"/> Appeal Communication to Board of Appeals and Interferences <input type="checkbox"/> Appeal Communication to TC (Appeal Notice, Brief, Reply Brief) <input type="checkbox"/> Proprietary Information <input type="checkbox"/> Status Letter <input checked="" type="checkbox"/> Other Enclosure(s) (please identify below): Return Postcard Nine (9) cited references
Remarks: The Commissioner is authorized to charge any additional fees to Deposit Account 20-1430.		

SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT

Firm Name	Townsend and Townsend and Crew LLP		
Signature			
Printed name	Chun-Pok Leung		
Date	October 24, 2005	Reg. No.	41,405

CERTIFICATE OF TRANSMISSION/MAILING

Express Mail Label: EV 529865174 US

I hereby certify that this correspondence is being deposited with the United States Postal Service with "Express Mail Post Office to Address" service under 37 CFR 1.10 on this date **October 24, 2005** and is addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on the date shown below.

Signature			
Typed or printed name	Joy Salvador	Date	October 24, 2005



PATENT
Attorney Docket No.: 16869B-077300US
Client Ref. No.: HAL264
(340300817US01)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of:

NAOKI WATANABE

Application No.: 10/767,247

Filed: January 28, 2004

For: METHOD AND APPARATUS
FOR COPYING AND BACKUP
IN STORAGE SYSTEMS

Customer No.: 20350

Examiner: Mano Padmanabhan

Technology Center/Art Unit: 2188

Confirmation No.: 4787

**PETITION TO MAKE SPECIAL FOR
NEW APPLICATION UNDER M.P.E.P.
§ 708.02, VIII & 37 C.F.R. § 1.102(d)**

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This is a petition to make special the above-identified application under MPEP § 708.02, VIII & 37 C.F.R. § 1.102(d). The application has not received any examination by an Examiner.

(a) The Commissioner is authorized to charge the petition fee of \$130 under 37 C.F.R. § 1.17(i) and any other fees associated with this paper to Deposit Account 20-1430.

(b) All the claims are believed to be directed to a single invention. If the Office determines that all the claims presented are not obviously directed to a single invention, then Applicants will make an election without traverse as a prerequisite to the grant of special status.

10/28/2005 TBESHAH1 00000018 201430 10767247
01 FC:1464 130.00 DA

(c) Pre-examination searches were made of U.S. issued patents, including a classification search and a key word search. The classification search was conducted on or around June 15, 2005 covering Class 707 (subclasses 200-202 and 204), Class 711 (subclasses 111-114, 147, 153, 161, 162, 165, 170, 173, and 203-209), and Class 714 (subclasses 1, 5, 8, 42, 47, 100, 723, 769, and 770), by a professional search firm, Mattingly, Stanger, Malur & Brundidge, P.C. The key word search was performed on the USPTO full-text database including published U.S. patent applications. A search for foreign art was also conducted using the European Patent Office's ESPACENET database and Japanese patent database.

(d) The following references, copies of which are attached herewith, are deemed most closely related to the subject matter encompassed by the claims:

- (1) U.S. Patent No. 5,287,363;
- (2) U.S. Patent No. 5,761,411;
- (3) U.S. Patent No. 5,815,650;
- (4) U.S. Patent No. 6,223,300 B1;
- (5) U.S. Patent No. 6,327,679 B1;
- (6) U.S. Patent No. 6,363,462 B1;
- (7) U.S. Patent No. 6,516,425 B1;
- (8) U.S. Patent Publication No. 2002/0059539 A1; and
- (9) U.S. Patent Publication No. 2002/0174295 A1.

(e) Set forth below is a detailed discussion of references which points out with particularity how the claimed subject matter is distinguishable over the references.

A. Claimed Embodiments of the Present Invention

The claimed embodiments relate to storage systems, and in particular to storage system management in which failure boundaries are taken into consideration when assigning storage volumes. A technique is provided for controlling a storage system in which

primary storage volumes and replication storage volumes are present. A boundary of a potential failure of the primary storage volumes and the replication storage volumes is determined, and using that boundary, replication storage volumes are assigned to assure that at least some of them are outside the failure boundary.

Independent claim 1 recites a method of controlling a storage system having primary storage volumes and replication storage volumes which replication storage volumes improve reliability of the storage system. The method comprises determining a boundary of a potential failure of the primary storage volumes and the replication storage volumes; and using the determined boundary to assign replication storage volumes to assure that at least some of the replication storage volumes are outside the failure boundary.

Independent claim 10 recites a storage system comprising a set of primary storage volumes; a set of replication storage volumes for improving reliability of the storage system; a memory for storing information regarding at least one boundary of a potential failure of the primary storage volumes and the replication storage volumes; and a controller coupled to the memory for assigning replication storage volumes to assure that at least some of the replication storage volumes are outside the failure boundary.

One of the benefits that may be derived is that it avoids the undesirable circumstance in which a physical failure may not only impact the primary volume, but also the replication volumes, where the storage controller will not consider the physical layout when it creates a replication pattern.

B. Discussion of the References

1. U.S. Patent No. 5,287,363

The patent to Wolf et al., US 5287363, discloses a method and apparatus for use in a computer having peripheral data storage, to monitor data transfers to detect defective or marginally defective storage areas. The method and apparatus are based on a memory-resident program component and a transient program component which cooperate. The resident component is in the computer's interrupt chain, to be invoked during both the initiation and conclusion of storage operations. A non memory resident part of the system of this invention retrieves data logged by the memory resident part, assembles the logged data

into a history of the medium per unit area, e.g., per sector, and analyzes the historical data to determine whether there are any areas of the storage medium that are unable to store data with integrity or are tending to lose the ability to store data with integrity. If any areas are determined to be bad or to be marginally bad as a result of the analysis performed by this part, the non memory resident part relocates to other areas the data stored in these areas, if possible, and marks original areas bad according to the procedures of the operating system under which the non-memory resident part is running. See column 3, lines 40-53; column ; column 17, lines 4-38.

Wolf et al. discloses determining whether storage areas are defective, but not determining a boundary of a potential failure of primary storage volumes and replication storage volumes. Wolf et al. discloses marking original areas bad and relocating the data if any areas are determined to be bad, but not using the boundary to assign replication storage volumes. More specifically, Wolf et al. fails to teach determining or storing a boundary of a potential failure of the primary storage volumes and the replication storage volumes; and using the boundary to assign replication storage volumes to assure that at least some of the replication storage volumes are outside the failure boundary, as recited in independent claims 1 and 10.

2. U.S. Patent No. 5,761,411

The patent to Teague et al., US 5761411, discloses a method for performing disk fault prediction operations. A layered block device driver for accessing a storage device coupled to a computer system having a platform on which a disk fault prediction application operates. The prediction of an imminent failure of the storage device is determined by selecting a threshold value for the selected attribute and, upon determining that the monitored attribute has crossed the selected threshold value, predicting the imminent failure of the storage device. In alternate aspects thereof, the selected attribute may be a performance attribute, an error rate attribute or an event count attribute. For the performance attribute, imminent failure is predicted upon detection of a degradation in the performance attribute such that the monitored performance attribute has a speed less than the threshold performance speed for the selected performance attribute. Figures 7A-B show the method of predicting an imminent failure of an ATA disk drive or other storage device. Continuous monitoring of the storage device and updating of attributes for which occurrences of events are detected is

maintained until the timer times out, thereby indicating that the time period between successive saves of the drive attributes to the storage device has elapsed. See column 2, line 63-column 3, line 14; column 9, line 35 to column 10, line 65.

Teague et al. discloses a method of predicting an imminent failure of a storage device, but not predicting a boundary of a potential failure of primary storage volumes and replication storage volumes and then using the boundary to assign replication storage volumes. More specifically, Teague et al. fails to teach determining or storing a boundary of a potential failure of the primary storage volumes and the replication storage volumes; and using the boundary to assign replication storage volumes to assure that at least some of the replication storage volumes are outside the failure boundary, as recited in independent claims 1 and 10.

3. U.S. Patent No. 5,815,650

The patent to Apperley et al., US 5815650, discloses an efficient system for predicting and processing storage subsystem failure. Predictive failure analysis of a storage subsystem is efficiently conducted and data quickly recovered from a failed Read operation. This may be implemented in a storage system 100 including a host 102 coupled to a supervising processor 104 that couples to a parity-equipped RAID storage subsystem 101 having multiple HDAs 108-113 each including an HDA controller 115 and at least one storage medium 127. In one embodiment, when an HDA experiences an error during a Read attempt, the HDA transmits a recovery alert signal to the supervising processor; then, the processor and HDA begin remote and local recovery processes in parallel. The first process to complete provides the data to the host, and the second process is aborted. See column 4, lines 40-47. In another embodiment, an HDA's PFA operations are restricted to idle times of the HDA. See column 5, lines 20-26. A different embodiment limits HDA performance of PFA to times when the processor is conducting data reconstruction. See column 5, line 63 to column 6, line 8. Another embodiment monitors HDA errors at the supervisory processor level, initiating an HDA's PFA operations when errors at that HDA have a certain characteristic, such as a predetermined frequency of occurrence. See column 6, lines 36-42.

Apperley et al. discloses monitoring HDA errors and data recovery, but not using a boundary of a potential failure of primary storage volumes and replication storage

volumes using the boundary to assign replication storage volumes. More specifically, Apperley et al. fails to teach determining or storing a boundary of a potential failure of the primary storage volumes and the replication storage volumes; and using the boundary to assign replication storage volumes to assure that at least some of the replication storage volumes are outside the failure boundary, as recited in independent claims 1 and 10.

4. U.S. Patent No. 6,223,300 B1

The patent to Gotoh, US 6223300, discloses a disk array apparatus which, without stopping its host system, reconstructs data areas of its configured storage devices by determining optimum array parameters in keeping with disk access trends. The apparatus 1 includes an access log storing means 5, 6 for storing access logs to the log storage areas; an access log analyzing means 7 for analyzing trends of access to the storage devices based on the access logs; an array parameter judging and changing means 8 for judging the validity of array parameters based on the result of the analyses and for setting optimum array parameters anew as needed; and a reconstructing means 9 for reconstructing the data areas as per the changed array parameters. Two methods of analyzing the access logs are disclosed. The first analyzing method involves analyzing data in the access logs on a time series basis. The second analyzing method involves sorting and analyzing the data in the access logs in terms of access sizes. The access logs may also be analyzed by statistical techniques. See column 5, lines 9-25.

Gotoh discloses analyzing access logs and judging and changing optimum array parameters. It does not determine a boundary of a potential failure of primary storage volumes and replication storage volumes, and does not use the boundary to assign replication storage volumes. More specifically, Gotoh fails to teach determining or storing a boundary of a potential failure of the primary storage volumes and the replication storage volumes; and using the boundary to assign replication storage volumes to assure that at least some of the replication storage volumes are outside the failure boundary, as recited in independent claims 1 and 10.

5. U.S. Patent No. 6,327,679 B1

The patent to Russell, US 6327679, discloses relocating unreliable disk sectors when encountering disk drive read errors with notification to user when data is bad. When a

read to a disk returns an error for a sector previously identified as good, the read is retried a predetermined number of times to attempt to recover the sector data. If the sector is unrecoverable, the failed sector is relocated and an "unusable" bit associated with the replacement sector is set to indicate that the data is bad. A defect map table, which may be maintained by storage media and/or the operating system for host system, contains entries for each LBN where an error has been detected. Until an unrecoverable sector is identified for storage media, defect map table will contain no entries. As unrecoverable sectors are identified over the life of storage media, entries are added to defect map table. When an unrecoverable sector is identified, the failed sector is mapped within defect map table to a replacement sector previously reserved by the operating system for host system. See column 4, lines 4-19. When an unrecoverable sector is identified during operation, the sector is remapped to a reserved spare or replacement sector. See column 4, lines 20-34.

Russell discloses identifying, mapping, and remapping unrecoverable sectors. It does not disclose determining a boundary of a potential failure of primary storage volumes and replication storage volumes, and using the boundary to assign replication storage volumes. More specifically, Russell fails to teach determining or storing a boundary of a potential failure of the primary storage volumes and the replication storage volumes; and using the boundary to assign replication storage volumes to assure that at least some of the replication storage volumes are outside the failure boundary, as recited in independent claims 1 and 10.

6. U.S. Patent No. 6,363,462 B1

The patent to Bergsten, US 6363462, discloses a storage controller providing automatic retention and deletion of synchronous back-up data. A network comprises at least one host processing system 2, a number of storage controllers 3-1 to 3-M, each coupled to one of a plurality of storage arrays 4-1 to 4-M, each storage array including at least one mass storage device. Each storage controller may be coupled to at least one host processing system and to at least one other storage controller to control access of the host processing systems to the mass storage devices. As shown in the routine for performing a standard read operation of FIG. 14, the standard checks routine is performed in step 1402. If any of the standards checks fail, then in step 1403 the storage controller identifies the appropriate logical device

based on the host interface and virtual block member received from the host. FIG. 19 illustrates a routine for performing automatic detection and correction of data errors.

Bergsten discloses detection and correction of data errors. It does not disclose determining a boundary of a potential failure of primary storage volumes and replication storage volumes, and using the boundary to assign replication storage volumes. More specifically, Bergsten fails to teach determining or storing a boundary of a potential failure of the primary storage volumes and the replication storage volumes; and using the boundary to assign replication storage volumes to assure that at least some of the replication storage volumes are outside the failure boundary, as recited in independent claims 1 and 10.

7. U.S. Patent No. 6,516,425 B1

The patent to Belhadj et al., US 6516425, discloses a method of managing data in a hierarchical data storage system employing data redundancy schemes includes prioritizing a data rebuild based on a most vulnerable data redundancy scheme in the storage system. Prioritizing the data rebuild includes enabling a rebuild of the most vulnerable data redundancy scheme prior to enabling a rebuild of any other data redundancy scheme in the system. The most vulnerable data redundancy scheme is determined by comparing a probability of losing data that can be prevented by a rebuild for each data redundancy scheme with respect to the potential for one or more next storage device failures in the data storage system. For instance, a most vulnerable redundancy scheme (i.e., RAID level) is defined as a redundancy scheme having the highest probability for potential data loss that can be prevented by a rebuild in the array 15 (relative to any other RAID levels within its redundancy group and relative to any other RAID levels within any other redundancy group in the array, all with respect to the potential for one or more next storage device failures in the array). Where a same redundancy scheme (i.e., RAID level) is employed in two or more separate redundancy groups, then a most vulnerable redundancy scheme may be defined relative to a given redundancy group. See column 8, lines 30-41.

Belhadj et al. discloses prioritizing data rebuild based on the most vulnerable data redundancy scheme. It does not disclose determining a boundary of a potential failure of primary storage volumes and replication storage volumes, and using the boundary to assign replication storage volumes. More specifically, Belhadj et al. fails to teach determining or

storing a boundary of a potential failure of the primary storage volumes and the replication storage volumes; and using the boundary to assign replication storage volumes to assure that at least some of the replication storage volumes are outside the failure boundary, as recited in independent claims 1 and 10.

8. U.S. Patent Publication No. 2002/0059539 A1

The published patent application to Anderson, US 20020059539, discloses a hybrid data storage and reconstruction system and method for a data storage device. The invention is drawn to a hybrid data reconstruction system and method for a data storage device. Data is selectively stored according to one of two or more redundancy schemes such that critical data is stored according to a scheme which has a higher degree of redundancy. The storage space of a storage medium includes a number of objects. Associated with each object is a set of attributes. In one example, an access control attribute is provided which is set by a Set Attribute method and provides means by which access to a particular object is controlled. By changing the version number of the access control attribute, certain requesters 116-120 can be denied or given, access to the particular object. See paragraph [0071]. The various attributes may include a clustering object that is an attribute which indicates whether the particular object should desirably be located near another object in the storage system. A cloning attribute indicates whether the particular object was created by copying another object in the storage system. A group of size attributes define the size characteristics of the particular object. A group of time attributes indicates when the object was created, the last time data in the object was modified, and the last time an attribute was modified in the object. See paragraphs [0072]-[0073].

Anderson discloses storing critical data according to a scheme which has a higher degree of redundancy. It does not teach determining a boundary of a potential failure of primary storage volumes and replication storage volumes, and using the boundary to assign replication storage volumes. More specifically, Anderson fails to teach determining or storing a boundary of a potential failure of the primary storage volumes and the replication storage volumes; and using the boundary to assign replication storage volumes to assure that at least some of the replication storage volumes are outside the failure boundary, as recited in independent claims 1 and 10.

9. U.S. Patent Publication No. 2002/0174295 A1

The published patent application to Ulrich et al., US 20020174295, discloses an enhanced file system with failure tolerance. The enhanced file system includes a plurality of disk drives for storing parity groups, each parity group comprising storage blocks. The storage blocks include one or more data blocks and a parity block associated with the data blocks. Each of the storage blocks is stored on a separate disk drive such that no two storage blocks from a given parity group reside on the same disk drive. A recovery module dynamically recovers data lost when a disk drive becomes unavailable. FIG. 30 illustrates a fault recovery mechanism 700 used by the distributed file storage system 100 to maintain data consistency and integrity when a data fault occurs. Data faults are characterized by corruption or loss of data or information stored in one or more logical blocks 2330 of the array 140. Each data fault can be further characterized as a catastrophic event, where an entire disk 2305 fails requiring all data on the failed disk to be reconstructed. Alternatively, the data fault can be characterized as a localized event, where the disk 2305 maintains operability but one or more physical disk sectors or logical blocks become corrupted or damaged. In either instance of the data fault, the distributed file storage system 100 uses a fault-tolerant restoration process to maintain data integrity. See paragraph [452].

Ulrich et al. discloses a fault recovery mechanism for maintaining data consistency and integrity in a distributed file storage system. It does not disclose determining a boundary of a potential failure of primary storage volumes and replication storage volumes, and using the boundary to assign replication storage volumes. More specifically, Ulrich et al. fails to teach determining or storing a boundary of a potential failure of the primary storage volumes and the replication storage volumes; and using the boundary to assign replication storage volumes to assure that at least some of the replication storage volumes are outside the failure boundary, as recited in independent claims 1 and 10.

Appl. No. 10/767,247
Petition to Make Special

PATENT

(f) In view of this petition, the Examiner is respectfully requested to issue
a first Office Action at an early date.

Respectfully submitted,



Chun-Pok Leung
Reg. No. 41,405

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834
Tel: 650-326-2400
Fax: 415-576-0300
Attachments
RL:rl
60601649 v1